Embracing Evolution: A Call for Body-Control Co-Design in Embodied Humanoid Robot

Anonymous Author(s)
Affiliation
Address
email

Abstract

Humanoid robots, as general-purpose physical agents, must integrate both intelligent control and adaptive morphology to operate effectively in diverse real-world environments. While recent research has focused primarily on optimizing control policies for fixed robot structures, this position paper argues for *"evolving both control strategies and humanoid robots' physical structure under a co-design mechanism"*. Inspired by biological evolution, this approach enables robots to iteratively adapt both their form and behavior to optimize performance within task-specific and resource-constrained contexts. Despite its promise, co-design in humanoid robotics remains a relatively underexplored domain, raising fundamental questions about its feasibility and necessity in achieving true embodied intelligence. To address these challenges, we propose practical co-design methodologies grounded in strategic exploration, Sim2Real transfer, and meta-policy learning. We further argue for the essential role of co-design by analyzing it from methodological, application-driven, and community-oriented perspectives. Striving for guiding and inspiring future studies, we present open research questions, spanning from short-term innovations to long-term goals. This work positions co-design as a cornerstone for developing the next generation of intelligent and adaptable humanoid agents.

1 Introduction

As an emerging research area, Embodied AI posits that intelligence stems from an agent's ability to actively explore, interact with, and learn from its environment in a continuous and dynamic manner. Within this learning paradigm, recent studies have developed various robot control models based on deep neural network backbones, enabling scalability across diverse tasks and environments [1, 2, 3, 4].

In studies of embodied robot agents, their skills are closely tied to the physical form. For example, robot arms, grippers, and dexterous hands are commonly employed for manipulation tasks such as grasping, placing, and assembling objects [5, 6]. Similarly, wheeled robots, bipedal robots, and quadrupedal robots are designed for locomotion tasks, including walking, climbing, and navigation [7]. To develop general-purpose robots, recent studies have focused on humanoid robots. Equipped with dual arms, legged body, and advanced sensors, humanoid robots are well-suited for a wide range of mobile locomotion tasks, enabling them to seamlessly handle everyday tasks [8, 9, 10, 11].

In recent years, the development of humanoid robots has primarily centered on control policy design, typically built upon predefined physical structures. These robotic designs are often the result of manual engineering and domain-specific heuristics, which are rarely subject to optimization within the embodied humanoid system. However, embodied intelligence is not solely determined by control performance, but is also fundamentally grounded in agents' physical structure [12]. For instance, in natural systems, organisms evolve their body morphology to adapt to changing environmental conditions. Similarly, embodied agents should incorporate evolutionary mechanisms to adapt to task requirements and environmental dynamics.

An effective method of realizing such evolutionary mechanisms is the robotic co-design problem, which seeks to jointly optimize both the control policy and the morphological design of robotic

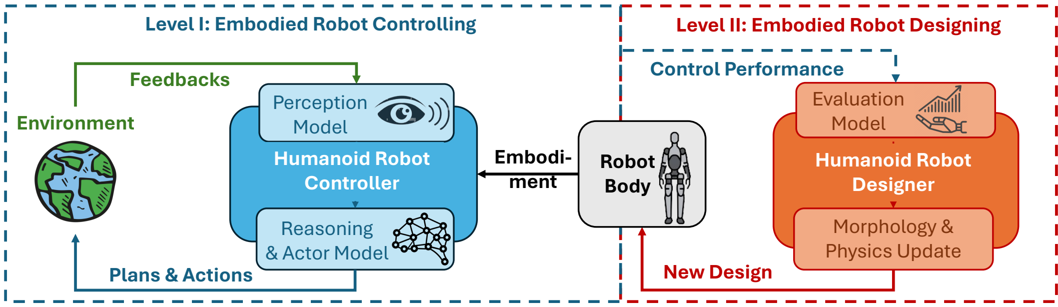

Figure 1: The co-design framework for humanoid robots, which can be formulated as a bi-level optimization problem, consisting of two interconnected phases: 1) learning the control policy for the humanoid robot, and 2) designing the robot's physical structure (Section 3.1).

systems [13]. While prior studies have explored co-design in quadruped robots [14, 15], soft robots [16], bi-piedal robots [17, 18] and modular robots [19, 20] (see Table 1), its extension to more advanced humanoid robots and connection to embodied intelligence remains largely unexplored. It remains unclear how to efficiently discover the optimal design of a generalist humanoid robot capable of performing a variety of tasks. More importantly, the necessity of addressing such co-design problems in the development of embodied humanoid robots has yet to be fully established.

This article provides a principled formulation of the humanoid co-design problem, emphasizing that **evolving physical structure is both feasible and essential for realizing embodied intelligence in humanoid robots**. Specifically, we formulate the humanoid co-design problem as a bi-level optimization. Such an optimizer can be integrated into the reasoning–acting architecture of an advanced controlling model, enabling an embodied humanoid robot to exhibit dexterity, mobility, perception, and intelligence.

Beyond the proposed formulation, we investigate an alternative perspective for realizing embodied humanoid robots based on predefined and manually specified designs. We analyze why such paradigms prevail in recent humanoid robotics research and examine the potential challenges of adopting co-design, particularly regarding algorithmic complexity, physical evaluation, and design scalability. To address these challenges, we introduce advancements in learning-based solvers, such as strategic robot structure exploration, the Sim2Real learning paradigm, and meta control policy, highlighting the feasibility of evolving humanoid robot architectures.

To understand the necessity of humanoid robot co-design, we investigate its unique advantages in facilitating robot morphology optimization, real-world task adaptation, and cross-disciplinary collaboration, examined from the perspectives of methodology, application, and community. To realize these key advantages, we identify open questions within the humanoid robot co-design problem, highlighting those that may be tractable with current methodologies in the short term, as well as those that may depend on long-term advances in emerging research domains.

2 Embodied Humanoid Robots

2.1 Architecture of Humanoid Robot

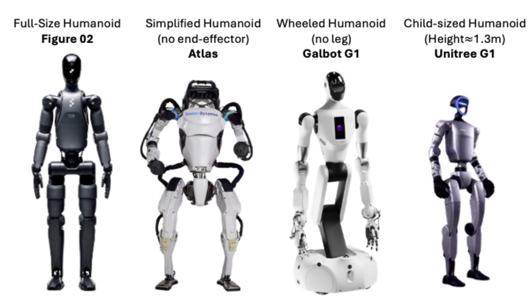

Figure 2: Examples of humanoid robots with varying physical structures are shown from left to right: full-sized, simplified, wheeled, and child-sized humanoid robots.

Humanoid robots are a specialized type of physical robot designed to replicate human-like functionality [21]. An ideal humanoid robot often has leggy designs in its *lower body*, featuring a *bi-pedal structure* that enables finishing locomotion tasks like walking, running, and maintaining balance. The *upper body* includes *dual arms* equipped with dexterous hands as end-effectors, allowing the robot to perform complex tasks that require precise manipulation and human-like hand movements. Their *sensor systems* commonly provide both *proprioception* and *exteroception*. Proprioception ensures internal body awareness by monitoring joint positions, angular velocity, and pose estimation, while exteroception enables perception of external states, such as LiDAR sweeps and RGB-D data.

While such designs are ideal, building and fine-tuning a humanoid robot's architecture typically demands substantial effort and resources. This often necessitates certain simplifications, such as 1) omitting end effectors in the dual arms, 2) constructing child-sized robots instead of full-sized ones, 3) utilizing a wheeled base for the lower body, and 4) downplaying the amounts and quality of sensors. Figure 2 provides illustrative examples of humanoid models.

Desideratas for Humanoid Robot. Given their human-like body structure, a fundamental requirement for humanoid robots is the ability to operate seamlessly within human environments. This enables them to collaborate closely with humans or take on dangerous or physically demanding tasks. During operation, the robot should exhibit natural behavior, adhering closely to human behavioral norms, even when performing long-term tasks across varying environments. These desiderata demand a control model with strong generalizability and adaptability, which traditional optimization-based control methods often struggle to achieve. Essentially, fulfilling this requirement calls for the development of embodied intelligence in humanoid robots.

2.2 Embodied Intelligence in Humanoid Robot

Unlike traditional approaches that rely on passively learning from fixed datasets, Embodied Artificial Intelligence (E-AI) requires agents to actively explore, interact with, and learn from their environment in a continuous and dynamic manner. Specifically, to enable the learning of an embodied humanoid robot, it often requires the robot to have four key abilities, including: 1) **Dexterity**: the ability to manipulate various objects with precision, delicacy, and intricacy. 2) **Mobility**: the capability to move and navigate through environments with different terrains and conditions. 3) **Perception**: the skill to gather, interpret, and understand environmental information from sensors. 4) **Intelligence**: the ability to process information, reason about sub-goals related to a given task, and adapt effectively to diverse tasks and environments.

These capabilities reflect not just the functionality of a humanoid robot in specific tasks like locomotion or manipulation but also emphasize generalization to a wide range of real-world scenarios, thereby advancing toward zero-shot deployment of humanoid robots for realistic applications.

2.3 Implementation for Embodied Humanoid Robot.

To achieve the above abilities, recent studies have implemented embodied humanoid robots using a two-layer architecture consisting of a high-level reasoning model and a low-level action model [22, 23, 24]. This hierarchical design is inspired by the functional organization of the human brain, where the cerebrum is responsible for logical reasoning and decision-making, while the cerebellum governs fine-grained motor control and coordination. These models are introduced as follows:

Humanoid Robot Reasoning Model. The reasoning model is typically implemented as a large-scale robotics foundation model designed to perform logical reasoning over the necessary steps for a robot to complete given tasks. The model takes as input natural language instructions describing the task, along with perceptual data, primarily visual signals collected from the environment. Its outputs consist of high-level plans to guide robotic execution. These may include sub-task descriptions and intermediate goals, as well as more detailed robotic outputs such as motion trajectories, grasp poses, and contact points (e.g., affordances). To learn this reasoning model, The training of the robotic reasoning model is typically achieved by fine-tuning a pre-existing Vision-Language Model (VLM) using refined robotic operation data, which is either collected through teleoperation or generated using synthetic data engines.

Humanoid Robot Action Model. The action model takes the outputs from the reasoning model (either as explicit data or implicit latent variables) as goals $g \in \mathcal{G}$ and predicts the corresponding control signals, such as joint angles or torques for the robot's joints at each time step. To learn action policies, previous methods commonly formulate the learning environment as a (partially observable) Markov Decision Process (MDP). $\mathcal{M} = (\mathcal{S}, \mathcal{A}, \mathcal{O}, P_{\mathcal{T}}, r, \mu_0, \gamma)$, where: 1) Within the state space $\mathcal{S}$, a state $s \in \mathcal{S}$ records the complete environmental information and the robot's internal states. 2) $\mathcal{A}$ denotes the action space, and action $a \in \mathcal{A}$ denotes the angles or torques at joints of the humanoid robot. 3) $o \in \mathcal{O}$ denotes the observations obtained from the robot's sensors, encompassing both proprioceptive inputs that reflect the humanoid's internal state and exteroceptive inputs that capture information about the external environment. 4) r denotes the reward functions, which typically consist of penalty, regularization, and task rewards. In particular, the magnitude of the task reward should closely reflect how well the robot accomplishes the given goal g. 5) $P_{\mathcal{T}} \in \Delta^{\mathcal{S}}_{\mathcal{S}\times\mathcal{A}}$ denotes the transition function as a mapping from state-action pairs to a distribution of future states. 6) $\mu_0 \in \Delta^{\mathcal{S}}$ denotes the initial state distribution. 7) $\gamma \in (0, 1]$ denotes the discounting factor.

Under this MDP, the humanoid action model can be represented as a meta policy $\pi \in \Delta^{\mathcal{A}}_{\mathcal{S}\times\mathcal{G}}$ that can scale to diverse goals $g\mathcal{G}$ under different environmental states $s \in \mathcal{S}$. During training, the goal is to maximize the expected cumulative discounted rewards:

$$\max_{\pi\in\Pi} \mathcal{J}(\pi, \mathcal{M}) = \max_{\pi\in\Pi} \mathbb{E}_{\mu_0,p_{\mathcal{T}},\pi}\left[\sum_{t=0}^{\infty} \gamma^t r(s_t, a_t, g)\right] \tag{1}$$

3 Position Proposal and Alternative Views

In the following sections, we introduce a co-design framework that jointly considers control policies and the evolution of humanoid morphology. Additionally, we present an alternative perspective: embodied intelligence should be grounded in predefined humanoid structures without evolution.

3.1 A New Perspective: Co-Designing Control and Evolution Policies

In natural environments, animals exhibit remarkable embodied intelligence, leveraging their evolved morphologies to learn and perform complex tasks [12]. Inspired by this, we argue that evolutionary principles should play an essential role in the development of embodied humanoid robots. While prior research commonly focused on perception, reasoning, and control within fixed robot structures, our position is that the robot's physical form itself should also be subject to optimization as a core component of its design. The simultaneous optimization of a humanoid robot's action model [1]) and physical components can be formulated as a co-design problem, integrating both control and morphology in the design process.

As illustrated in Figure 2, the robot co-design problem requires the joint optimization of both control policies and physical robot modules to maximize overall performance, while adhering to resource constraints such as cost [13]. When extending this framework to a learning-based setting, the co-design task typically involves a forward pass for training the control policy and a backward pass for updating the robot's physical parameters.

Specifically, during the forward process, we use the RL algorithm to optimize the goal-aware policy function by maximizing $\mathcal{J}(\pi, \mathcal{M}_\psi)$ in the objective (1), where the configuration of this learning environment (i.e., MDP) depends on a specific physical physical robot structure, denoted by $\psi \in \Psi$. Based on the policy performance, we conduct an inverse update on the robot's structure, which necessitates formulating humanoid robot co-design as a bi-level optimization problem. Moreover, the design often incorporates system-level constraints, which define strict requirements for the desired system behavior (e.g., resembling human behavior) or impose limitations on the resources (e.g., costs of robot modules) [13]. The optimization problem can be described as:

$$\max_{\psi\in\Psi} \max_{\pi\in\Pi} \mathcal{J}(\pi, \mathcal{M}_\psi) \text{ s.t. } f_c(\psi) \leq \epsilon \tag{2}$$

3.2 Alternative Views: Intelligence Arises from Fixed Humanoid Robots Structure

While the co-design approach provides a significantly broader design space for enhancing the performance of humanoid robots, most recent studies, spanning manipulation [25, 26, 27], locomotion [28, 29, 30, 31], and human motion imitation (i.e., teleoperation) [32, 8, 10, 11, 9, 33], continue to focus on controlling fixed, pre-defined humanoid platforms, without considering structural modifications to the robot itself. Even with the recent surge in exploring embodied intelligence in humanoid robots across a wide range of everyday tasks [34], *robotic co-design, as an effective technique in robotics research, remains largely underexplored* in this context [35]. This trend reflects a prevailing assumption that *"The predefined and fixed physical structures are sufficient for supporting the development embodied humanoid robots"*.

This perspective essentially treats the robot's structure as a fixed component of the environment's dynamics, which can be estimated and adapted to, but not actively optimized. Conceptually, this assumption is prevalent in the RL literature [36], which serves as a foundational algorithm for

[1]Since compare to robot acting, task reasoning is typically less dependent on the robot's detailed physical structure and updating a VLM is often costly, the reasoning model is generally not updated alongside changes to the robot's morphology.

learning-based control in humanoid robots and inherently shapes subsequent research directions. More importantly, in practice, there are significant challenges associated with co-designing humanoid robots, further reinforcing the reliance on manually designed, fixed-structure humanoid platforms.

1) Complexity of the Co-design Problem. In addition to learning control policy, the co-design problem incorporates an second-level optimization loop for refining the robot's physical design [13, 12, 37, 15, 17, 38]. For humanoid robots, this process becomes especially challenging due to their complex upper and lower body structures, which often involve varying configurations of motors, joints, sensors, and body components. Consequently, exploring the high-dimensional design space and identifying optimal configurations demands substantial computational resources. In many cases, due to the intricate interdependencies between the design and control parameters, the bi-level optimization in the co-design problem may struggle to converge. Without additional restrictions, the optimal solution may not be uniquely identifiable or even computationally tractable.

2) Difficulties in Physical Evaluation. To evaluate the optimality of a humanoid robot structure, it is crucial to deploy the robot in task-specific scenarios and assess how effectively it can adapt its control model to complete those tasks. This process often requires modifying certain components of the robot based on the proposed design. However, unlike simpler robotic systems, humanoid robots have highly complex and interdependent structures [39, 21, 40]. Modifying one part frequently leads to changes in the robot's overall physical configuration. For instance, adjusting the length of the thighs affects the robot's weight distribution and center of mass. These changes, in turn, influence both kinematic properties (e.g., motions and velocities) and dynamic characteristics (such as inertia and gravity models). The technical challenges in physical reconfiguration limit the feasibility of iterative structural design and evaluation in real-world applications.

3) Limited Scalability Across Tasks. Robotic co-design typically aims to enhance performance for specific tasks [13]. For instance, [17] optimized the leg length of a bipedal robot to achieve maximum walking velocity. Similarly, [15, 41] explored the joint optimization of mechanical structures and control policies to improve the locomotion capabilities of quadruped robots. However, embodied humanoid robots, with physical structures resembling those of humans, are designed to generalize across a wide range of tasks and environments within human workspaces. This requires multi-dimensional capabilities, including dexterity, mobility, perception, and intelligence (Section 2.2). The task-specific optimization frameworks commonly used in traditional robotic co-design cannot be directly applied to the inherently cross-task nature of humanoid robots. This lack of task scalability limits the overall utility of co-designed systems, particularly when targeting general-purpose humanoid platforms intended for reuse across diverse applications.

4 Feasibility of Co-Designing Humanoid Robot

To address the inherent challenges of humanoid robot co-design, this section investigates its feasibility by proposing a set of potential solutions. In particular, we introduce three key strategies: strategic exploration, the Sim2Real paradigm, and meta-policy learning. These approaches are aimed at tackling critical issues in co-design, including the complexity of joint design and control, the challenges of physical evaluation, and the limited scalability across diverse tasks. Moreover, by leveraging recent advances in control algorithms, simulated environments, foundation models, and decision-making policies, these strategies establish a robust foundation for the development of next-generation co-design algorithms for embodied humanoid systems.

4.1 Strategic Exploration under Constrained Design Space

In the co-design literature, genetic algorithms take a critical role in modifying robot structures via crossover, mutation, and replacement operations [12, 14, 15, 17, 38]. During this process, their underlying exploration mechanisms are inherently random. This randomness results in unstructured exploration, lacking informative priors or guidance toward designs that are more likely to yield higher-performing robots. This challenge becomes especially pronounced in unstructured and high-dimensional design spaces, such as those encountered in the development of humanoid robots, which induces computational burden (see alternative views in Section 3.2).

To address this limitation, *strategic exploration* has emerged as a promising approach for accelerating the search process. In the RL literature, a variety of algorithms have been developed to promote more efficient exploration [42]. For instance, recent work has proposed provably efficient strategies based on Bayesian updates [43, 44], the Upper Confidence Bound (UCB) [45], and uncertaintydriven heuristics [46, 47], achieving significantly lower regret bounds compared to purely random exploration. Motivated by these advances, we propose formulating the physical physical robot structure process as an MDP, which enables the application of strategic exploration techniques to more effectively explore the design space of humanoid robots.

In addition to accelerating exploration, another important strategy for ensuring computational tractability and design identifiability is to *constrain the design space* with the following methods: 1) Rather than modifying the entire robot structure, *the exploration can be limited to a few key modules or components*. For instance, [18, 41] focused on optimizing the thigh and shank lengths of humanoid robots. Similarly, [37] investigated the configuration of various motors and the inertial parameters of robot links. Other studies, such as [48, 15], examined the parameters of parallel elastic knee joints. 2) From a theoretical perspective, a key approach to ensuring the convergence of bi-level optimization is to *bound the range of the design parameters*. By utilizing a compact design space and bounding the objective function in Equation 2, the Extreme Value Theorem (EVT) [49] guarantees the existence of a maximum. In this context, given a continuous objective function, the bi-level optimization can converge to an optimal (but not necessarily unique) solution (π^*, ψ^*).

By strategically exploring a compact and bounded design space focused on key modules of the humanoid robot, the computational burden of humanoid robot co-design can be significantly reduced. This approach substantially alleviates the complexity of the bi-level optimization process (Section 3.2).

4.2 A Sim2Real Paradigm for Evaluation and Deployment

Guiding structural updates of humanoid robots typically requires evaluation signals from the current design. Traditional robot co-design studies often involve building physical hardware and assessing its performance in real-world tasks [20, 18, 15, 48]. However, as discussed in alternative perspectives (Section 3.2), this real-world design and evaluation approach is not directly applicable to humanoid robots due to the complexity of their physical structures.

Instead of relying on real-world evaluations, an alternative and effective approach is to conduct both the design and evaluation of robots within simulated environments. For example, studies such as [50, 51] and [15, 41, 17] utilize simulators like MuJoCo and Isaac Gym to learn control policies and evaluate the task performance of various physical robot structures. While simulation reduces the cost and complexity of building physical hardware, the Simulation-to-Reality (Sim2Real) gap remains a major challenge: a robot that performs well in simulation may not exhibit the same level of performance in the real world. Additionally, for humanoid robots with soft components, accurately simulating their morphology, especially for contact-rich interactions, remains difficult [52].

As a result, developing a feasible Sim2Real paradigm for humanoid co-design requires effective strategies to bridge the gap between simulation and reality. Key approaches include: 1) Extending *domain randomization and domain adaptation* techniques to humanoid control environments, which remains a critical direction for future research [53]. 2) *Developing simulation platforms with high photorealism and physical fidelity* to minimize the impact of the Sim2Real gap. For instance, recently developed simulation environments and engines such as RoboCasa [34], Isaac Lab [54], MuJoCo-Playground [55], Genesis [56], and ManiSkill [57] have demonstrated promising capabilities in accurately modeling real-world scenes and physical interactions. 3) Beyond explicitly modeling objects, scenarios, and physical laws, recent work has introduced the concept of the *World Function Model* [58, 59]. This paradigm treats simulation as a regression task, where the model predicts future states of the environment in response to perturbations (i.e., actions). This offers a more flexible and data-driven alternative to traditional simulators.

4.3 Controlling Humanoid Robot via Meta Policy

As discussed in Section 3.2, traditional approaches to robotic design and control have primarily focused on optimizing simple robots to complete specific tasks [13]. In contrast, embodied humanoid robots, as complex, legged, dual-arm systems equipped with numerous sensors, are designed to learn generalizable policies that can be applied across a wide range of tasks and environments. This fundamental difference renders previously proposed robot co-design methods not directly transferable to the problem of controlling embodied humanoid robots.

To enable the designed robot to perform effectively across multiple tasks, a critical approach is to learn a meta-policy π that allows the humanoid robot to solve a wide range of everyday human tasks involving both manipulation and locomotion [34]. In the context of humanoid control, the meta-policy can be modeled as a goal-conditioned policy $\pi : \mathcal{S} \times \mathcal{G} \to \mathcal{A}$, where a goal $g \in \mathcal{G}$

may represent various forms of task specification, such as commands [31, 11, 29], target poses [32, 8], affordances [60, 61], or more generally, natural language descriptions of tasks [22, 24, 23]. Recent advances in generalizable decision models, such as decision transformers [62] and diffusion policies [1], provide effective backbone architectures for implementing such meta-policies, enabling flexible and scalable control across a wide range of goal specifications.

To integrate robotic co-design into policy learning, goals can be randomly sampled from a predefined goal pool during training. The robot's policy is then conditioned on each goal and adapted according to a proposed design configuration ψ. Similarly, during evaluation, the effectiveness of a given physical robot structure can be assessed by measuring how well it facilitates the learning of a policy that successfully achieves a diverse set of goals across varying environments.

5 Necessity of Co-Designing Humanoid Robot

In the previous section, we explored the feasibility of jointly modeling the robot's physical structure and its control policy, outlining key strategies that make such co-design tractable. In this section, we go a step further and argue for the necessity of co-design in the development of embodied humanoid robots from the perspective of metrology, application, and community.

5.1 Methodology: Principled Optimization of Robot Morphology

While significant progress has been achieved using pre-designed humanoid robots in both locomotion and manipulation tasks, there remains a lack of principled methods for evaluating the optimality of these designs. In practice, to determine robot morphology in Figure 2, *it commonly relies on engineers' intuition and experience, rather than through systematic optimization*. Such practice is often inefficient, since the robot's physical design is independent of the training of its planning models and control policies. The full capabilities and limitations of a given design often only become apparent when other research groups attempt to tackle more complex locomotion or manipulation tasks, revealing shortcomings that were not initially evident. These insights are used retrospectively to inform the development of robots, which typically progresses slowly due to the lack of systematic design methodologies. For example, in the initial design of the Unitree H1 robot, the limited degrees of freedom (DoFs) in its arms significantly constrained its ability to perform manipulation tasks that require rich interactions with objects. Recognizing this limitation, the developers addressed it in the subsequent version, Unitree H1-2 [63], by increasing the number of DoFs in each arm from 4 to 7. However, this design revision took over a year to implement.

In addition to improving the efficiency of robot development, co-design can significantly enhance its *efficacy*. From an algorithmic standpoint, allowing the robot's design ψ to vary enables the learning algorithm to explore a larger joint search space over both morphology and control. This expanded space allows for the discovery of design-policy pairs that maximize overall performance:

$$\max_{\psi \in \Psi} \max_{\pi \in \Pi} \mathcal{J}(\pi, \mathcal{M}_{\psi}) \geq \max_{\pi \in \Pi} \mathcal{J}(\pi, \mathcal{M}_{\psi'}), \quad \forall \psi' \in \Psi \tag{3}$$

Here, $\mathcal{J}$ denotes the expected return of policy π in the environment defined by the morphology $\mathcal{M}_{\psi}$. This inequality highlights the potential performance gains from jointly optimizing both the robot's design and its control policy. Without principled optimization, manually identifying the optimal design $\psi^* \in \Psi$ is highly challenging, particularly for humanoid robots expected to perform diverse tasks and adapt to complex, real-world environments encountered in everyday scenarios. *To optimize both the efficacy and efficiency of discovering effective humanoid body designs, we argue that humanoid robot co-design is essential*.

5.2 Application: Adaptive Body Shaping for Real-World Tasks

In practice, the specific requirements for a humanoid robot's capabilities vary depending on the deployment environment of each real-world application. For example, in industrial settings, humanoid robots are often tasked with manipulating a variety of objects for relocation, rearrangement, or assembly. In such scenarios, dexterity becomes a critical factor, as robots must precisely and efficiently handle components of different shapes, sizes, and material properties [64, 30, 31, 65, 28]. In contrast, when deployed as patrol robots in environments such as university campuses or public facilities, humanoid robots must robustly navigate to different locations under diverse terrains (e.g., slopes, stairs) and environmental disturbances (e.g., dynamic obstacles or weather conditions). In these applications, robustness and mobility become the key performance criteria [66, 67, 26, 25]. Most importantly, *there exists a fundamental trade-off between dexterity and mobility* in the design of a humanoid robot's body structure. Achieving dexterous manipulation typically requires highly flexible arms (with many DoFs) and delicate actuators. However, these design choices often result in increased weight and a higher center of mass, which can influence the robot's balance and reduce the efficacy of locomotion tasks.

A similar trade-off observed in humanoid robots can also be found in human physiology. For example, the body composition of boxers and runners differs significantly in terms of muscle distribution and weight allocation. These athletes often dedicate substantial time to optimizing their bodies to enhance the specific skills required in their respective sports. Just as athletes undergo intensive training camps to simultaneously develop both their physical form and technical skills, the co-design process of a humanoid robot's body and control policy can be viewed as a training camp for humanoid robots. By continuously adapting both morphology and behavior across diverse applications, we can dynamically tailor robot structures that are best suited for the target tasks and desired applications. *We argue that the co-design is essential for the humanoid robot to adapt specifically to its target applications.*

5.3 Community: Fostering Cross-Disciplinary Collaboration

Co-designing the control model and body structure of an embodied humanoid robot is fundamentally a multidisciplinary research topic. It encompasses: 1) *Machine learning expertise* for processing multi-modal sensory inputs, learning control policies, and performing high-level planning; 2) *Robotics design principles* for modeling and optimizing the robot's dynamics and kinematics; and 3) *Mechanical engineering knowledge* for the manufacturing of structural components and the integration of hardware systems. Each of these topics represents a significant research field with its own dedicated communities and research groups.

Traditionally, these research communities have evolved and been explored independently. For instance, machine learning is closely aligned with data-driven AI approaches, often emphasizing theoretical and methodological advancements in software systems. In contrast, robotics design and mechanical engineering emphasize the physical realization of hardware systems. However, when it comes to humanoid robot co-design, it fundamentally relies on the joint optimization of these domains, due to their deep and intricate interdependencies. Its advancement demands interdisciplinary collaboration, and the development of unified frameworks capable of optimizing control algorithms, morphological design, and hardware implementation in a coherent and efficient manner.

In recent years, cross-disciplinary collaboration, particularly under the AI+"X" paradigm, has played a vital role in advancing scientific and technological breakthroughs. For example, AlphaFold [68] exemplifies the synergy between deep neural networks and structural biology, revolutionizing protein structure prediction. OpenAI Five [69] and AlphaStar [70] integrate deep RL with the gaming and entertainment industry, pushing the boundaries of AI in complex, multi-agent environments. Similarly, Med-PaLM [71] bridges LLMs with medical knowledge, enabling AI-assisted healthcare solutions. These successes highlight the transformative potential of combining AI with domain-specific expertise. In this context, we argue that *humanoid robot co-design is essential for its unique capacity to foster cross-disciplinary collaboration across AI, robotics, and engineering.*

6 Open Questions in Humanoid Robot Co-Design

To encourage further exploration of humanoid robot co-design, we propose a set of open questions that may be addressed in both the short and long term.

Open Questions in Short-Term. We present open research questions that could be effectively tackled by leveraging emerging techniques and models.

1) *Efficient Representation for Robot Design.* Deriving concise and informative representations of data has been a key factor in the success of modern machine learning. In robotic learning tasks, recent studies have explored efficient representations for various types of multi-modal data, including language [72], 2D images [73], 3D shapes such as point clouds [74] and scenes [75, 76], and tactile information [77]. However, in most of these studies, the robot structure is fixed, and there has been limited exploration of efficient representations for robot morphology. This lack of focus significantly limits the efficiency of learning and adaptation in tasks involving adaptive robot design, particularly those that rely on updating deep neural networks.

2) *Benchmarking Robot Co-Design.* In robot co-design, Sim2Real training plays a crucial role by allowing the performance of co-design algorithms to be evaluated in simulated environments before deployment on physical hardware (Section 4.2). However, unlike other robotic manipulations and

location tasks with rich benchmarks [78, 79, 5, 80, 7], humanoid robot co-design, as a relatively new research area, lacks commonly applied benchmarks. Instead, prior studies often customize their tasks and environments according to specific goals, making it difficult to assess how well these algorithms generalize to other settings. While these case-specific studies provide valuable insights for physical robot structure, their applicability to other embodied humanoid robots, especially in diverse tasks and environments, remains uncertain. As a result, establishing a standardized benchmark for humanoid robot co-design emerges as an important and timely objective for the field, especially with the availability of simulation platforms such as Isaac Gym [81], MuJoCo [82], and Genesis [56].

3) *Design-Aware Policy Optimization.* In addition to learning meta-policies that can adapt to different tasks and environments, an important open question is how to develop design-aware policies $\pi_\psi : \mathcal{S} \times \mathcal{G} \times \Psi \to \mathcal{A}$ [38]. Such policies are designed to generalize effectively across a range of different robot morphologies, enabling adaptive action control even when the physical structure ψ changes. When a modification in the robot's body occurs, the policy can still perform reasonably well and, with minimal fine-tuning, adapt to the new structure. In this way, the policy can effectively serve as a morphology-aware controller, reducing the need for retraining from scratch whenever structural changes are introduced. This capability is crucial for co-design frameworks, where iterative updates to both control and morphology are expected. Ultimately, robust generalization across morphologies not only accelerates the Sim2Real transfer process but also enhances the practicality and scalability of humanoid robot deployment in dynamic, real-world environments.

Open Questions in Long-Run. We introduce promising robotic co-design research topics that depend on the advancement of other emerging areas, which are actively being studied but have yet to yield effective solutions.

1) *Co-Designs with World Models.* While learning-based co-design heavily relies on simulated environments, the simulators typically use manually specified semantics, rules, and physical laws, resulting in a non-negligible gap between simulation and the real world. To address this issue, recent studies [59, 58] have proposed building World Function Models (WFMs) that learn dynamics directly from real-world data. Inspired by the success of foundation models, WFMs are data-driven systems that automatically learn real-world physics and dynamics based on actions, without relying on human-designed assumptions. By jointly optimizing both the control policy and the robot morphology structure using WFMs, the gap between simulation and real-world application can be significantly reduced. However, developing reliable WFMs remains a challenging and long-term objective. As such, co-design based on WFMs is expected to be a major goal for future research.

2) *Co-Design Planning and Reasoning*. Current co-design studies primarily focus on the joint optimization of the action model and the robot structure. Although planning and reasoning models are integral components of embodied humanoid control systems, they are typically not subject to optimization during co-design and are instead used as pre-trained models. A major reason for this is the heavy computational burden associated with inferring and updating these large-scale Vision-Language Models (VLMs). In contrast, action models (i.e., policies) are relatively smaller, making their optimization more computationally manageable. As a result, a promising future direction is exploring the joint optimization of both reasoning and action models, so as the cover the entire process of. Achieving this will require the development of highly efficient inference and learning techniques (e.g., the use of Mixture-of-Experts (MoE) architectures), which remains an important and active area of ongoing research.

7 Conclusion

This paper advocates for a body-control co-design paradigm in humanoid robotics, emphasizing the joint optimization of both control strategies and physical morphology. Inspired by principles of biological evolution, we argue that co-design is essential for achieving embodied intelligence, enabling humanoid robots to adapt more effectively to diverse, dynamic real-world tasks. We demonstrate the feasibility of this approach through strategic exploration, Sim2Real transfer, and meta-policy learning, and highlight its necessity across methodological, application-driven, and interdisciplinary perspectives. By integrating co-design into the development pipeline, we can embrace the potential for more robust, generalizable, and intelligent humanoid systems. To guide future research, we outline key open questions, ranging from representation learning, benchmarking, and design-aware controlling to long-term integration with world models and reasoning systems. We position co-design as a foundational approach for developing intelligent, adaptable, and general-purpose humanoid robots capable of thriving in complex real-world environments.

References

[1] Tony Z. Zhao, Vikash Kumar, Sergey Levine, and Chelsea Finn. Learning fine-grained bimanual manipulation with low-cost hardware. In *Robotics: Science and Systems, RSS*, 2023.

[2] Ji Lin, Hongxu Yin, Wei Ping, Pavlo Molchanov, Mohammad Shoeybi, and Song Han. VILA: on pre-training for visual language models. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR*, pages 26679–26689. IEEE, 2024.

[3] Songming Liu, Lingxuan Wu, Bangguo Li, Hengkai Tan, Huayu Chen, Zhengyi Wang, Ke Xu, Hang Su, and Jun Zhu. RDT-1b: a diffusion foundation model for bimanual manipulation. In *International Conference on Learning Representations, ICLR*, 2025.

[4] Dibya Ghosh, Homer Rich Walke, Karl Pertsch, Kevin Black, Oier Mees, Sudeep Dasari, Joey Hejna, Tobias Kreiman, Charles Xu, Jianlan Luo, You Liang Tan, Lawrence Yunliang Chen, Quan Vuong, Ted Xiao, Pannag R. Sanketi, Dorsa Sadigh, Chelsea Finn, and Sergey Levine. Octo: An open-source generalist robot policy. In *Robotics: Science and Systems, RSS*, 2024.

[5] Jianlan Luo, Charles Xu, Fangchen Liu, Liam Tan, Zipeng Lin, Jeffrey Wu, Pieter Abbeel, and Sergey Levine. Fmb: a functional manipulation benchmark for generalizable robotic learning. *The International Journal of Robotics Research, IJRR*, 2024.

[6] Jianlan Luo, Charles Xu, Jeffrey Wu, and Sergey Levine. Precise and dexterous robotic manipulation via human-in-the-loop reinforcement learning. *arXiv preprint arXiv:2410.21845*, 2024.

[7] Zhuoqun Xu, Yang Liu, Xiaoqi Li, Jiyao Zhang, and Hao Dong. Human-centered in-building embodied delivery benchmark. *arXiv preprint arXiv:2406.17898*, 2024.

[8] Zipeng Fu, Qingqing Zhao, Qi Wu, Gordon Wetzstein, and Chelsea Finn. Humanplus: Humanoid shadowing and imitation from humans. In *8th Annual Conference on Robot Learning*, 2024.

[9] Xuxin Cheng, Yandong Ji, Junming Chen, Ruihan Yang, Ge Yang, and Xiaolong Wang. Expressive whole-body control for humanoid robots. *arXiv preprint arXiv:2402.16796*, 2024.

[10] Tairan He, Zhengyi Luo, Xialin He, Wenli Xiao, Chong Zhang, Weinan Zhang, Kris Kitani, Changliu Liu, and Guanya Shi. Omnih2o: Universal and dexterous human-to-humanoid whole-body teleoperation and learning. *arXiv preprint arXiv:2406.08858*, 2024.

[11] Tairan He, Wenli Xiao, Toru Lin, Zhengyi Luo, Zhenjia Xu, Zhenyu Jiang, Jan Kautz, Changliu Liu, Guanya Shi, Xiaolong Wang, Linxi Fan, and Yuke Zhu. HOVER: versatile neural whole-body controller for humanoid robots. *arXiv preprint arXiv:2410.21229*, 2024.

[12] Agrim Gupta, Silvio Savarese, Surya Ganguli, and Li Fei-Fei. Embodied intelligence via learning and evolution. *Nature communications*, 12(1):5721, 2021.

[13] Luca Carlone and Carlo Pinciroli. Robot co-design: Beyond the monotone case. In *International Conference on Robotics and Automation, ICRA*, pages 3024–3030, 2019.

[14] Álvaro Belmonte-Baeza, Joonho Lee, Giorgio Valsecchi, and Marco Hutter. Meta reinforcement learning for optimal design of legged robots. *IEEE Robotics and Automation Letters*, 7(4):12134–12141, 2022.

[15] Filip Bjelonic, Joonho Lee, Philip Arm, Dhionis V. Sako, Davide Tateo, Jan Peters, and Marco Hutter. Learning-based design and control for quadrupedal robots with parallel-elastic actuators. *IEEE Robotics and Automation Letters*, 8(3):1611–1618, 2023.

[16] Yuxing Wang, Shuang Wu, Tiantian Zhang, Yongzhe Chang, Haobo Fu, Qiang Fu, and Xueqian Wang. Preco: Enhancing generalization in co-design of modular soft robots via brain-body pre-training. In Jie Tan, Marc Toussaint, and Kourosh Darvish, editors, *Conference on Robot Learning, CoRL*, volume 229 of *Proceedings of Machine Learning Research*, pages 478–498, 2023.

[17] Yi Cheng, Chenxi Han, Yuheng Min, Houde Liu, Linqi Ye, and Hang Liu. Structural optimization of lightweight bipedal robot via SERL. In *IEEE/RSJ International Conference on Intelligent Robots and Systems,IROS*, pages 9066–9073, 2024.

[18] Adrian B. Ghansah, Jeeseop Kim, Maegan Tucker, and Aaron D. Ames. Humanoid robot co-design: Coupling hardware design with gait generation via hybrid zero dynamics. In *IEEE Conference on Decision and Control, CDC*, pages 1879–1885, 2023.

[19] Allan Zhao, Jie Xu, Mina Konakovic-Lukovic, Josephine Hughes, Andrew Spielberg, Daniela Rus, and Wojciech Matusik. Robogrammar: graph grammar for terrain-optimized robot design. *ACM Transactions on Graphics*, 39(6):188:1–188:16, 2020.

[20] Julian Whitman, Raunaq M. Bhirangi, Matthew J. Travers, and Howie Choset. Modular robot design synthesis with deep reinforcement learning. In *AAAI Conference on Artificial Intelligence*, pages 10418–10425, 2020.

[21] Saeed Saeedvand, Masoumeh Jafari, Hadi S. Aghdasi, and Jacky Baltes. A comprehensive survey on humanoid robot development. *Knowledge Engineering Review*, 34:e20, 2019.

[22] Gemini Robotics Team, Saminda Abeyruwan, Joshua Ainslie, Jean-Baptiste Alayrac, Montserrat Gonzalez Arenas, Travis Armstrong, Ashwin Balakrishna, Robert Baruch, Maria Bauza, Michiel Blokzijl, et al. Gemini robotics: Bringing ai into the physical world. *arXiv preprint arXiv:2503.20020*, 2025.

[23] Haoqi Yuan, Yu Bai, Yuhui Fu, Bohan Zhou, Yicheng Feng, Xinrun Xu, Yi Zhan, Börje F Karlsson, and Zongqing Lu. Being-0: A humanoid robotic agent with vision-language models and modular skills. *arXiv preprint arXiv:2503.12533*, 2025.

[24] Figure AI. Helix: A vision-language-action model for generalist humanoid control, February 2025. Accessed: 2025-04-15.

[25] Yanjie Ze, Zixuan Chen, Wenhao Wang, Tianyi Chen, Xialin He, Ying Yuan, Xue Bin Peng, and Jiajun Wu. Generalizable humanoid manipulation with improved 3d diffusion policies. *arXiv preprint arXiv:2410.10803*, 2024.

[26] Jinhan Li, Yifeng Zhu, Yuqi Xie, Zhenyu Jiang, Mingyo Seo, Georgios Pavlakos, and Yuke Zhu. OKAMI: Teaching humanoid robots manipulation skills through single video imitation. In *8th Annual Conference on Robot Learning, CoRL*, 2024.

[27] Xuxin Cheng, Jialong Li, Shiqi Yang, Ge Yang, and Xiaolong Wang. Open-television: Teleoperation with immersive active visual feedback. *CoRR*, abs/2407.01512, 2024.

[28] Zixuan Chen, Xialin He, Yen-Jen Wang, Qiayuan Liao, Yanjie Ze, Zhongyu Li, S Shankar Sastry, Jiajun Wu, Koushil Sreenath, Saurabh Gupta, et al. Learning smooth humanoid locomotion through lipschitz-constrained policies. *arXiv preprint arXiv:2410.11825*, 2024.

[29] Ziwen Zhuang, Shenzhe Yao, and Hang Zhao. Humanoid parkour learning. In *8th Annual Conference on Robot Learning, CoRL*, 2024.

[30] Junfeng Long, Junli Ren, Moji Shi, Zirui Wang, Tao Huang, Ping Luo, and Jiangmiao Pang. Learning humanoid locomotion with perceptive internal model. *arXiv preprint arXiv:2411.14386*, 2024.

[31] Xinyang Gu, Yen-Jen Wang, Xiang Zhu, Chengming Shi, Yanjiang Guo, Yichen Liu, and Jianyu Chen. Advancing humanoid locomotion: Mastering challenging terrains with denoising world model learning. *arXiv preprint arXiv:2408.14472*, 2024.

[32] Tairan He, Zhengyi Luo, Wenli Xiao, Chong Zhang, Kris Kitani, Changliu Liu, and Guanya Shi. Learning human-to-humanoid real-time whole-body teleoperation. *arXiv preprint arXiv:2403.04436*, 2024.

[33] Mazeyu Ji, Xuanbin Peng, Fangchen Liu, Jialong Li, Ge Yang, Xuxin Cheng, and Xiaolong Wang. Exbody2: Advanced expressive humanoid whole-body control. *arXiv preprint arXiv:2412.13196*, 2024.

[34] Soroush Nasiriany, Abhiram Maddukuri, Lance Zhang, Adeet Parikh, Aaron Lo, Abhishek Joshi, Ajay Mandlekar, and Yuke Zhu. Robocasa: Large-scale simulation of everyday tasks for generalist robots. *arXiv preprint arXiv:2406.02523*, 2024.

[35] Yuchuang Tong, Haotian Liu, and Zhengtao Zhang. Advancements in humanoid robots: A comprehensive review and future prospects. *IEEE/CAA Journal of Automatica Sinica*, 11(2):301–328, 2024.

[36] Richard S Sutton. Reinforcement learning: An introduction. *A Bradford Book*, 2018.

[37] Carlotta Sartore, Lorenzo Rapetti, Fabio Bergonti, Stefano Dafarra, Silvio Traversaro, and Daniele Pucci. Codesign of humanoid robots for ergonomic collaboration with multiple humans via genetic algorithms and nonlinear optimization. In *IEEE-RAS International Conference on Humanoid Robots, Humanoids*, pages 1–8, 2023.

[38] Gabriele Fadini and Stelian Coros. Making use of design-aware policy optimization in legged-robotics co-design. In *[CoRL 2024] Morphology-Aware Policy and Design Learning Workshop (MAPoDeL)*.

[39] Dusko Katic and Miomir Vukobratovic. Survey of intelligent control techniques for humanoid robots. *Journal of Intelligent and Robotic Systems*, 37(2):117–141, 2003.

[40] Kourosh Darvish, Luigi Penco, Joao Ramos, Rafael Cisneros, Jerry E. Pratt, Eiichi Yoshida, Serena Ivaldi, and Daniele Pucci. Teleoperation of humanoid robots: A survey. *IEEE Trans. Robotics*, 39(3):1706–1727, 2023.

[41] Ci Chen, Jiyu Yu, Haojian Lu, Hongbo Gao, Rong Xiong, and Yue Wang. Pretraining-finetuning framework for efficient co-design: A case study on quadruped robot parkour. *arXiv preprint arXiv:2407.06770*, 2024.

[42] Pawel Ladosz, Lilian Weng, Minwoo Kim, and Hyondong Oh. Exploration in deep reinforcement learning: A survey. *Information Fusion*, 85:1–22, 2022.

[43] Sreejith Balakrishnan, Quoc Phong Nguyen, Bryan Kian Hsiang Low, and Harold Soh. Efficient exploration of reward functions in inverse reinforcement learning via bayesian optimization. In *Advances in Neural Information Processing Systems, NeurIPS*, 2020.

[44] Yishay Mansour, Aleksandrs Slivkins, Vasilis Syrgkanis, and Zhiwei Steven Wu. Bayesian exploration: Incentivizing exploration in bayesian games. *Operations Research*, 70(2):1105–1127, 2022.

[45] Xiuyuan Lu and Benjamin Van Roy. Information-theoretic confidence bounds for reinforcement learning. In *Advances in Neural Information Processing Systems, NeurIPS*, pages 2458–2466, 2019.

[46] Qi Cai, Zhuoran Yang, Chi Jin, and Zhaoran Wang. Provably efficient exploration in policy optimization. In *International Conference on Machine Learning, ICML*, volume 119, pages 1283–1294, 2020.

[47] Chi Jin, Zhuoran Yang, Zhaoran Wang, and Michael I. Jordan. Provably efficient reinforcement learning with linear function approximation. In *Conference on Learning Theory, COLT*, volume 125, pages 2137–2143, 2020.

[48] Gabriel Bravo-Palacios, He Li, and Patrick M Wensing. Engineering compliance in legged robots via robust co-design. *IEEE/ASME Transactions on Mechatronics*, 2024.

[49] Stephen P Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

[50] Charles B. Schaff, David Yunis, Ayan Chakrabarti, and Matthew R. Walter. Jointly learning to construct and control agents using deep reinforcement learning. In *International Conference on Robotics and Automation, ICRA*, pages 9798–9805. IEEE, 2019.

[51] Kevin Sebastian Luck, Heni Ben Amor, and Roberto Calandra. Data-efficient co-adaptation of morphology and behaviour with deep reinforcement learning. In *Annual Conference on Robot Learning, CoRL*, volume 100, pages 854–869. PMLR, 2019.

[52] Charles B. Schaff, Audrey Sedal, and Matthew R. Walter. Soft robots learn to crawl: Jointly optimizing design and control with sim-to-real transfer. In *Robotics: Science and Systems, RSS*, 2022.

[53] Garrett Wilson and Diane J Cook. A survey of unsupervised deep domain adaptation. *ACM Transactions on Intelligent Systems and Technology (TIST)*, 11(5):1–46, 2020.

[54] Mayank Mittal, Calvin Yu, Qinxi Yu, Jingzhou Liu, Nikita Rudin, David Hoeller, Jia Lin Yuan, Ritvik Singh, Yunrong Guo, Hammad Mazhar, Ajay Mandlekar, Buck Babich, Gavriel State, Marco Hutter, and Animesh Garg. Orbit: A unified simulation framework for interactive robot learning environments. *IEEE Robotics and Automation Letters*, 8(6):3740–3747, 2023.

[55] Kevin Zakka, Baruch Tabanpour, Qiayuan Liao, Mustafa Haiderbhai, Samuel Holt, Jing Yuan Luo, Arthur Allshire, Erik Frey, Koushil Sreenath, Lueder A. Kahrs, Carlo Sferrazza, Yuval Tassa, and Pieter Abbeel. Mujoco playground: An open-source framework for gpu-accelerated robot learning and sim-to-real transfer., 2025.

[56] Genesis Authors. Genesis: A universal and generative physics engine for robotics and beyond, December 2024.

[57] Stone Tao, Fanbo Xiang, Arth Shukla, Yuzhe Qin, Xander Hinrichsen, Xiaodi Yuan, Chen Bao, Xinsong Lin, Yulin Liu, Tse-kai Chan, et al. Maniskill3: Gpu parallelized robotics simulation and rendering for generalizable embodied ai. *arXiv preprint arXiv:2410.00425*, 2024.

[58] Wenyan Cong, Hanqing Zhu, Peihao Wang, Bangya Liu, Dejia Xu, Kevin Wang, David Z Pan, Yan Wang, Zhiwen Fan, and Zhangyang Wang. Can test-time scaling improve world foundation model? *arXiv preprint arXiv:2503.24320*, 2025.

[59] Niket Agarwal, Arslan Ali, Maciej Bala, Yogesh Balaji, Erik Barker, Tiffany Cai, Prithvijit Chattopadhyay, Yongxin Chen, Yin Cui, Yifan Ding, et al. Cosmos world foundation model platform for physical ai. *arXiv preprint arXiv:2501.03575*, 2025.

[60] Chenhao Lu, Xuxin Cheng, Jialong Li, Shiqi Yang, Mazeyu Ji, Chengjing Yuan, Ge Yang, Sha Yi, and Xiaolong Wang. Mobile-television: Predictive motion priors for humanoid whole-body control. *CoRR*, abs/2412.07773, 2024.

[61] Chong Zhang, Wenli Xiao, Tairan He, and Guanya Shi. Wococo: Learning whole-body humanoid control with sequential contacts. In *Annual Conference on Robot Learning*, 2024.

[62] Lili Chen, Kevin Lu, Aravind Rajeswaran, Kimin Lee, Aditya Grover, Michael Laskin, Pieter Abbeel, Aravind Srinivas, and Igor Mordatch. Decision transformer: Reinforcement learning via sequence modeling. In *Advances in Neural Information Processing Systems, NeurIPS*, pages 15084–15097, 2021.

[63] Unitree. Unitree h1 / h1-2 unitree's first universal humanoid robot, 2025.

[64] Scott Kuindersma, Robin Deits, Maurice F. Fallon, Andres Valenzuela, Hongkai Dai, Frank Permenter, Twan Koolen, Pat Marion, and Russ Tedrake. Optimization-based locomotion planning, estimation, and control design for the atlas humanoid robot. *Autonomous Robots*, 40(3):429–455, 2016.

[65] Qiang Zhang, Peter Cui, David Yan, Jingkai Sun, Yiqun Duan, Gang Han, Wen Zhao, Weining Zhang, Yijie Guo, Arthur Zhang, and Renjing Xu. Whole-body humanoid robot locomotion with human reference. In *IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS*, pages 11225–11231, 2024.

[66] Christian Smith, Yiannis Karayiannidis, Lazaros Nalpantidis, Xavi Gratal, Peng Qi, Dimos V. Dimarogonas, and Danica Kragic. Dual arm manipulation - A survey. *Robotics Auton. Syst.*, 60(10):1340–1353, 2012.

[67] Aude Billard and Danica Kragic. Trends and challenges in robot manipulation. *Science*, 364(6446):eaat8414, 2019.

[68] Josh Abramson, Jonas Adler, Jack Dunger, Richard Evans, Tim Green, Alexander Pritzel, Olaf Ronneberger, Lindsay Willmore, Andrew J Ballard, Joshua Bambrick, et al. Accurate structure prediction of biomolecular interactions with alphafold 3. *Nature*, 630(8016):493–500, 2024.

[69] Christopher Berner, Greg Brockman, Brooke Chan, Vicki Cheung, Przemysław Dębiak, Christy Dennison, David Farhi, Quirin Fischer, Shariq Hashme, Chris Hesse, et al. Dota 2 with large scale deep reinforcement learning. *arXiv preprint arXiv:1912.06680*, 2019.

[70] Oriol Vinyals, Igor Babuschkin, Wojciech M Czarnecki, Michaël Mathieu, Andrew Dudzik, Junyoung Chung, David H Choi, Richard Powell, Timo Ewalds, Petko Georgiev, et al. Grandmaster level in starcraft ii using multi-agent reinforcement learning. *nature*, 575(7782):350–354, 2019.

[71] Karan Singhal, Tao Tu, Juraj Gottweis, Rory Sayres, Ellery Wulczyn, Mohamed Amin, Le Hou, Kevin Clark, Stephen R Pfohl, Heather Cole-Lewis, et al. Toward expert-level medical question answering with large language models. *Nature Medicine*, pages 1–8, 2025.

[72] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013.

[73] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning, ICML*, pages 8748–8763, 2021.

[74] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *IEEE conference on Computer Vision and Pattern Recognition, CVPR*, pages 652–660, 2017.

[75] Ben Mildenhall, Pratul P Srinivasan, Matthew Tancik, Jonathan T Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM*, 65(1):99–106, 2021.

[76] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4), 2023.

[77] Jialiang Zhao, Yuxiang Ma, Lirui Wang, and Edward Adelson. Transferable tactile transformers for representation learning across diverse sensors and tasks. In *Annual Conference on Robot Learning, CoRL*, 2024.

[78] Stephen James, Zicong Ma, David Rovick Arrojo, and Andrew J. Davison. Rlbench: The robot learning benchmark & learning environment. *IEEE Robotics and Automation Letters*, 5(2):3019–3026, 2020.

[79] Chengshu Li, Ruohan Zhang, Josiah Wong, Cem Gokmen, Sanjana Srivastava, Roberto Martín-Martín, Chen Wang, Gabrael Levine, Michael Lingelbach, Jiankai Sun, Mona Anvari, Minjune Hwang, Manasi Sharma, Arman Aydin, Dhruva Bansal, Samuel Hunter, Kyu-Young Kim, Alan Lou, Caleb R Matthews, Ivan Villa-Renteria, Jerry Huayang Tang, Claire Tang, Fei Xia, Silvio Savarese, Hyowon Gweon, Karen Liu, Jiajun Wu, and Li Fei-Fei. BEHAVIOR-1k: A benchmark for embodied AI with 1,000 everyday activities and realistic simulation. In *6th Annual Conference on Robot Learning*, 2022.

[80] Carmelo Sferrazza, Dun-Ming Huang, Xingyu Lin, Youngwoon Lee, and Pieter Abbeel. Humanoidbench: Simulated humanoid benchmark for whole-body locomotion and manipulation. *arXiv preprint arXiv:2403.10506*, 2024.

[81] Viktor Makoviychuk, Lukasz Wawrzyniak, Yunrong Guo, Michelle Lu, Kier Storey, Miles Macklin, David Hoeller, Nikita Rudin, Arthur Allshire, Ankur Handa, and Gavriel State. Isaac gym: High performance GPU based physics simulation for robot learning. In *Neural Information Processing Systems Track on Datasets and Benchmarks 1, NeurIPS Datasets and Benchmarks*, 2021.

[82] Emanuel Todorov, Tom Erez, and Yuval Tassa. Mujoco: A physics engine for model-based control. In *IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS*, pages 5026–5033. IEEE, 2012.

[83] Sehoon Ha, Stelian Coros, Alexander Alspach, Joohyung Kim, and Katsu Yamane. Joint optimization of robot design and motion parameters using the implicit function theorem. In *Robotics: Science and Systems, RSS*, volume 13, pages 10–15607, 2017.

[84] Ruta Desai, Beichen Li, Ye Yuan, and Stelian Coros. Interactive co-design of form and function for legged robots using the adjoint method. *arXiv preprint arXiv:1801.00385*, 2018.

[85] Muhan Li, David Matthews, and Sam Kriegman. Reinforcement learning for freeform robot design. In *International Conference on Robotics and Automation, ICRA*, pages 8799–8806, 2024.

[86] Junru Song, Yang Yang, Huan Xiao, Wei Peng, Wen Yao, and Feifei Wang. LASer: Towards diversified and generalizable robot design with large language models. In *The Thirteenth International Conference on Learning Representations*, 2025.

[87] Jiaheng Hu, Julian Whitman, Matthew Travers, and Howie Choset. Modular robot design optimization with generative adversarial networks. In *International Conference on Robotics and Automation, ICRA*, pages 4282–4288, 2022.

[88] Donald Joseph Hejna III, Pieter Abbeel, and Lerrel Pinto. Task-agnostic morphology evolution. In *International Conference on Learning Representations, ICLR*, 2021.

[89] Deepak Pathak, Christopher Lu, Trevor Darrell, Phillip Isola, and Alexei A Efros. Learning to control self-assembling morphologies: a study of generalization via modularity. *Advances in Neural Information Processing Systems, NeurIPS*, 32, 2019.

[90] Tingwu Wang, Yuhao Zhou, Sanja Fidler, and Jimmy Ba. Neural graph evolution: Automatic robot design. In *International Conference on Learning Representations, ICLR*, 2019.

[91] Haofei Lu, Zhe Wu, Junliang Xing, Jianshu Li, Ruoyu Li, Zhe Li, and Yuanchun Shi. Bodygen: Advancing towards efficient embodiment co-design. In *The Thirteenth International Conference on Learning Representations, ICLR*, 2025.

[92] Andrew Spielberg, Brandon Araki, Cynthia Sung, Russ Tedrake, and Daniela Rus. Functional co-optimization of articulated robots. In *2017 IEEE International Conference on Robotics and Automation, ICRA*, pages 5035–5042. IEEE, 2017.

[93] Muhan Li, Lingji Kong, and Sam Kriegman. Generating freeform endoskeletal robots. In *The Thirteenth International Conference on Learning Representations, ICLR*, 2025.

[94] Xiaomeng Xu, Huy Ha, and Shuran Song. Dynamics-guided diffusion model for sensor-less robot manipulator design. In *Conference on Robot Learning, CoRL*, volume 270, pages 4446–4462, 2024.

A A Summary of Recent Studies in Robotic Co-Design

Table 1: The summary of recent studies and progress in robotic co-design.

Designing Method	Robot Type	Designing Parameters
Meta RL [14]	Quadrupedal Robot	Thigh and shank lengths; gear ratios of the actuators.
Bayesian Optimization [15]	Quadrupedal Robot	Parameters of parallel elastic knee joint.
ADMM [48]	Quadrupedal Robot	Parameters of parallel elastic actuation (PEA).
Bayesian Optimization [41]	Quadrupedal Robot	Thigh and shank lengths.
Implicit Function Theorem [83]	Quadrupedal Robot	Link length; actuator poses.
Adjoint Method [84]	Quadruped and Hexapod Robot	Link lengths; actuator poses; robot's width and length
Evolution RL [17]	Lightweight Bipedal Robot	Thigh and shin lengths.
HZD Optimization [18]	Bipedal Robot	Thigh and shin lengths.
PPO [85]	Modular Soft Robot	3D voxel-wise material assignments and spatial placement.
LLM-aided Evolution Search [86]	Modular Soft Robot	3D voxel-wise material assignments and spatial placement.
Model Order Reduction [52]	Modular Soft Robot	The combination of actuator placement and pressure regulators.
DQN [20]	Modular Manipulating Robot	The combination of different modules.
RoboGAN [87]	Modular Locomoting Robot	Module type assignment on fixed-topology graph.
Graph Neural Network [88]	Modular Locomoting Robot	Size and position of limbs, type and range of joints
Particle Swarm Optimization [51]	Modular Locomoting Robot	Leg segment lengths.
PPO [89]	Modular Locomoting Robot	The combination of limbs.
Neural Graph Evolution [90]	Modular Locomoting Robot	The combination of different modules.
PPO [91]	Modular Locomoting Robot	Limb length and size; joint rotation range and torque limit.
Quadratic Programming [92]	Ariticulated Robot	Geometry and inertia of links; torque limits of joints.
PPO [50]	Legged Locomoting Robot	Link length and mass.
Genetic Algorithm [37]	ErgoCub2 Humanoid Robot	Motor types and link inertial parameters.
CMA-ES [93]	Freedom Endoskeletal Robot	Limb length; soft and rigid radii.
DGDM [94]	Sensor-less Jaw Manipulator	Manipulator finger geometry (represented as Bézier curves).
Binary Programming [13]	Autonomous Racing Drone	The combination of different modules.

Table 1 summarizes recent studies in robotic co-design. We observe that most of these works focus on relatively simple robot types, such as modular and quadrupedal robots, with a limited number of design parameters considered. More importantly, the majority of these studies are tailored to specific tasks and environments. Extending these co-design methods to more complex humanoid robots operating across a diverse range of tasks and settings remains an important but largely unexplored challenge.

No checklist is needed for position paper track (See Position Paper Track FAQ).